

AMENDMENTS TO THE CLAIMS

1. (Currently amended) A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the method comprising:

maintaining a record of locations to which data are expected to be written on the primary storage subsystem by a host processor, as indicated by a predetermined prediction algorithm based on the locations to which the data have already been written;

receiving the data from the host processor at the primary storage subsystem to be written to a specified location on the first non-volatile storage media;

if the specified location is not included in the record, updating the record using the prediction algorithm so that the record includes both the specified location and one or more further locations that have not yet been specified by the host processor;

outputting an acknowledgment to the host processor to indicate that the data have been stored in the data storage system after receiving the data and after updating the record if the specified location was not included in the record prior to updating the record;

copying the data from the primary storage subsystem to the secondary storage subsystem so as to create a mirror on the secondary storage subsystem of the data received by the primary storage subsystem; ~~and~~

storing the data in the specified location on both the first and second non-volatile storage media; and

upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record, a portion of the data from the secondary storage subsystem to the primary storage subsystem for storage on the primary storage subsystem.

2. (Original) The method according to claim 1, wherein copying the data comprises transmitting the data between mutually-remote sites over a communication link between the sites.

3. (Canceled)

4. (Currently amended) The method according to ~~claim 3~~ claim 1, and comprising, upon occurrence of a failure in the primary storage subsystem, configuring the secondary storage subsystem to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

5. (Canceled)

6. (Original) The method according to claim 1, wherein maintaining and updating the record comprise marking respective bits in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

7. (Original) The method according to claim 1, wherein maintaining the record comprises storing the record on the first non-volatile storage media, and wherein updating the record comprises modifying the record that is stored on the first non-volatile storage media.

8. (Original) The method according to claim 7, wherein modifying the record comprises:

comparing the specified location to a copy of the record held in a volatile memory on the primary storage subsystem;

modifying the copy of the record so that at least the specified location is included in the copy of the record; and

destaging the modified copy of the record to the first non-volatile storage media.

9. (Original) The method according to claim 8, wherein the record is not modified on the first non-volatile

storage media responsively to receiving the data as long as the specified location to which the data are to be written is included in the record.

10. (Original) The method according to claim 7, wherein modifying the record comprises adding a plurality of locations, including the specified location, to the record.

11. (Original) The method according to claim 1, wherein updating the record comprises predicting one or more further locations to which the host processor is expected to write the data in a subsequent write operation, and adding the one or more further locations to the record.

12. (Original) The method according to claim 11, wherein predicting the one or more further locations comprises selecting a predetermined number of consecutive locations in proximity to the specified location.

13. (Original) The method according to claim 11, wherein maintaining the record comprises recording the locations to which the data are written using an object-based storage technique, and wherein predicting the one or more further locations comprises choosing the one or more further locations based on a logical connection between storage objects.

14. (Original) The method according to claim 1, wherein updating the record comprises removing one or more locations, other than the specified location, from the record, so as to limit a size of the record.

15. (Original) The method according to claim 14, wherein removing the one or more locations comprises receiving an acknowledgment from the secondary storage subsystem that the data have been stored in the one or more locations on the second non-volatile storage media, and removing the one or more locations from the record responsively to the acknowledgment.

16. (Previously presented) The method according to claim 14, wherein removing the one or more locations comprises identifying the locations at which the first and second non-volatile storage media contain identical data, and selecting for removal one of the identified locations that was least-recently added to the record.

17. (Currently amended) A data storage system, comprising:

a primary storage subsystem, which comprises first non-volatile storage media; and

a secondary storage subsystem, which comprises second non-volatile storage media,

wherein the primary storage subsystem is arranged to receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, which is arranged to store the data in the specified location on the second non-volatile storage media so as to create a mirror on the secondary storage subsystem of the data received by the primary storage subsystem, and

wherein the primary storage subsystem is arranged to maintain a record of locations to which data are expected to be written on the primary storage subsystem by the host processor, as indicated by a predetermined prediction algorithm based on the locations to which the data have already been written, and upon receiving the data from the host processor, to update the record using the prediction algorithm so that the record includes both the specified location and one or more further locations that have not yet been specified by the host processor if the specified location is not included in the record, and to output an acknowledgement to the host processor to indicate that the data have been stored in the data storage system after receiving the data and after

updating the record if the specified location was not included in the record prior to updating the record, and wherein upon recovery of the system from a failure of the primary storage subsystem, the secondary storage subsystem is arranged to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

18. (Original) The system according to claim 17, wherein the first and second non-volatile storage media are located at mutually-remote sites, and wherein at least one of the primary and secondary storage subsystems is arranged to transmit the data over a communication link between the sites.

19. (Canceled)

20. (Currently amended) The system according to ~~claim 19~~ claim 17, wherein upon occurrence of a failure in the primary storage subsystem, the secondary storage subsystem is configurable to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

21. (Canceled)

22. (Original) The system according to claim 17, wherein the record comprises a bitmap, and wherein the primary storage subsystem marks respective bits in the bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

23. (Original) The system according to claim 17, wherein the primary storage subsystem is arranged to store and update the record on the first non-volatile storage media.

24. (Original) The system according to claim 23, wherein the primary storage subsystem comprises a volatile memory

and is arranged to hold a copy of the record in the volatile memory, and to update the record by modifying the copy of the record, and destaging the modified copy of the record to the first non-volatile storage media.

25. (Original) The system according to claim 24, wherein the record is not modified on the first non-volatile storage media responsively to receiving the data as long as the specified location to which the data are to be written is included in the record.

26. (Original) The system according to claim 23, wherein the primary storage subsystem is arranged, when the specified location is not included in the record, to update the record in the first non-volatile storage media by adding a plurality of locations, including the specified location, to the record.

27. (Original) The system according to claim 17, wherein the primary storage subsystem is arranged, when the specified location is not included in the record, to predict one or more further locations to which the host processor is expected to write the data in a subsequent write operation, and to add both the specified location and the one or more further locations to the record.

28. (Original) The system according to claim 27, wherein the one or more further locations predicted by the primary storage subsystem comprise a predetermined number of consecutive locations in proximity to the specified location.

29. (Original) The system according to claim 27, wherein the primary storage subsystem is arranged to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

30. (Original) The system according to claim 17, wherein the primary storage subsystem is arranged, upon updating

the record, to remove one or more locations, other than the specified location, from the record, so as to limit a size of the record.

31. (Original) The system according to claim 30, wherein the secondary storage subsystem is arranged to transmit an acknowledgment to the primary storage subsystem indicating that the data have been stored in the one or more locations on the second non-volatile storage media, and wherein the primary storage subsystem is arranged to remove the one or more locations from the record responsively to the acknowledgment.

32. (Previously presented) The system according to claim 30, wherein the primary storage subsystem is arranged to identify the locations at which the first and second non-volatile storage media contain identical data, and to remove from the record one of the identified locations that was least-recently added to the record.

33. (Currently amended) A computer software product for use in a data storage system including primary and secondary storage subsystems, which include respective first and second control units and respective first and second non-volatile storage media, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first and second control units, cause the first control unit to receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, and cause the second control unit to store the data in the specified location on the second non-volatile storage media so as to create a mirror on the secondary storage subsystem of the data received by the primary storage subsystem,

wherein the instructions further cause the first control unit to maintain a record of locations to which data are expected to be written on the primary storage subsystem by the host processor, as indicated by a predetermined prediction algorithm based on the locations to which the data have already been written, and upon receiving the data from the host processor, to update the record using the prediction algorithm so that the record includes both the specified location and one or more further locations that have not yet been specified by the host processor if the specified location is not included in the record, and to output an acknowledgement to the host processor to indicate that the data have been stored in the data storage system after receiving the data and after updating the record if the specified location was not included in the record prior to updating the record, and

wherein upon recovery of the system from a failure of the primary storage subsystem, the instructions cause the second control unit to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

34. (Original) The product according to claim 33, wherein the first and second non-volatile storage media are located at mutually-remote sites, and wherein the instructions cause at least one of the first and second control units to transmit the data over a communication link between the sites.

35. (Canceled)

36. (Currently amended) The product according to ~~claim 35~~ claim 33, wherein the instructions cause the secondary storage subsystem, upon occurrence of a failure in the primary storage subsystem, to serve as the primary



storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

37. (Canceled)

38. (Original) The product according to claim 33, wherein the record comprises a bitmap, and wherein the instructions cause the first control unit to mark respective bits in the bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

39. (Original) The product according to claim 33, wherein the instructions cause the first control unit to store and update the record on the first non-volatile storage media.

40. (Original) The product according to claim 39, wherein the instructions cause the first control unit to hold a copy of the record in a volatile memory of the primary storage subsystem, and to update the record by modifying the copy of the record, and destaging the modified copy of the record to the first non-volatile storage media.

41. (Original) The product according to claim 40, wherein the instructions cause the first control unit not to modify the record on the first non-volatile storage media responsively to receiving the data as long as the specified location to which the data are to be written is included in the record.

42. (Original) The product according to claim 39, wherein the instructions cause the first control unit, when the specified location is not included in the record, to update the record in the first non-volatile storage media by adding a plurality of locations, including the specified location, to the record.

43. (Original) The product according to claim 33, wherein the instructions cause the first control unit,

when the specified location is not included in the record, to predict one or more further locations to which the host processor is expected to write the data in a subsequent write operation, and to add both the specified location and the one or more further locations to the record.

44. (Original) The product according to claim 43, wherein the one or more further locations predicted by the first control unit comprise a predetermined number of consecutive locations in proximity to the specified location.

45. (Original) The product according to claim 43, wherein the instructions cause the first control unit to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

46. (Original) The product according to claim 33, wherein the instructions cause the first control unit, upon updating the record, to remove one or more locations, other than the specified location, from the record, so as to limit a size of the record.

47. (Original) The product according to claim 46, wherein the instructions cause the second control unit to transmit an acknowledgment to the primary storage subsystem indicating that the data have been stored in the one or more locations on the second non-volatile storage media, and further cause the first control unit to remove the one or more locations from the record responsively to the acknowledgment.

48. (Previously presented) The product according to claim 46, wherein the instructions cause the first control unit to identify the locations at which the first and second non-volatile storage media contain identical

49266A2

data, and to remove from the record one of the identified locations that was least-recently added to the record.